Pacific Packet Radio Society
Post Office Box 51562
Palo Alto, CA 94303 USA

First Class

H.S. Magnuski KA6M
2019 Barbara Drive
Palo Alto , CA 94303

Digitized by the Internet Archive
in 2025 with funding from
Amateur Radio Digital Communications, Grant 151

https://archive.org/details/pprsnews-2-3

Pacific Packet Radio Society Newsletter

Volume 2 Number 3 March 1986

Pacific Packet Radio Society (PPRS)
Post Office Box 51562
Palo Alto, CA 94303

March 4th Meeting
7:30 PM at the Ampex Cafeteria, 411 Broadway, Redwood City

Effective Packet Radio BBS Usage

Tips and tricks on using the area's bulletin board systems effectively from their operators and users, including Roy, AA4RE, the operator of the Garlic Valley's W0RLI-clone BBS. If you have questions or suggestions concerning VHF BBS operation, this should be a good time to get them answered.
Also... report on the recent TAPR annual meeting, latest routing information, and of course news and recent developments in packet operation and regulation.

President's Letter

AJ6T>PPRS:

FCC AUTOCONTROL DOCKET CONTROVERSY CONTINUES

Shock waves from the FCC Report and Order on PR 85-105 continue to reverberate throughout the packet radio community. The R&O permits automatic operation of amateur stations in routine digital communications on VHF, but the FCC took the opportunity to emphasize that they would not change exisiting rules which prohibit automatic retransmission of third party traffic. This prohibition appears to make bulletin boards (and perhaps all digipeaters) illegal unless a control operator is present. Also, no HF automatic operation is allowed at all.

Amateur response has ranged from anger to apathy. Some are submitting Petitions for Reconsideration to the FCC. Some are afraid to ask any questions on the "let sleeping dogs lie" theory (however, it's too late for that attitude since the FCC could hardly make things any worse at this point). In an incredible head-in-the-sand move, the PPRS membership voted not to petition the FCC for reconsideration.

According to Perry Williams, W1UED, ARRL Secretary and Washington Area Coordinator, a few East Coast HF bulletin board stations have gone off line, but the majority remain in operation. Contrary to rumor, NO AMATEURS HAVE BEEN CITED YET BY THE FCC. A Commission monitoring station did telephone at least a few HF BBS operators on the East Coast, apparantly to ascertain whether they were unattended.

Paul Rinaldo, W4RI, QST Editor, has written an 11 page ARRL Petition for Reconsideration which points out the impracticality of real time monitoring by control operators at all points in an automatic packet switched network. The ARRL is seeking to place the responsibility for proper handling of third party traffic on the originating station, and they are trying to convince the FCC not to consider ham-to-ham messages via digipeaters to be third party traffic.

W1UED is optimistic about the future of packet radio. He thinks that the FCC was not attempting to frustrate amateurs, but to protect us from commercial interests inside and outside our own ranks. I am also guardedly optimistic about the prospects for packet radio. I think that the FCC will declare that amateur-to-amateur transmissions through an automatic network are not third party traffic (and so can be handled without the presence of control operators at each node), and my guess is that BBS operation will be legal also. This opens the door for a full blown nationwide automatic packet switching network on VHF. However, HF BBS and other automatic operation probably will remain illegal (largely because the ARRL is opposed to it) so HF gateways will disappear. Perhaps the Commission will be more ameanable to HF networking next year after we show them some responsible VHF operation without third party abuses.

The near hysteria exhibited by some operators is unwarranted, and the rumor and disinformation mill here in the Bay Area does not reflect well on ham radio.

Walter E Miller, AJ6T 20Feb86

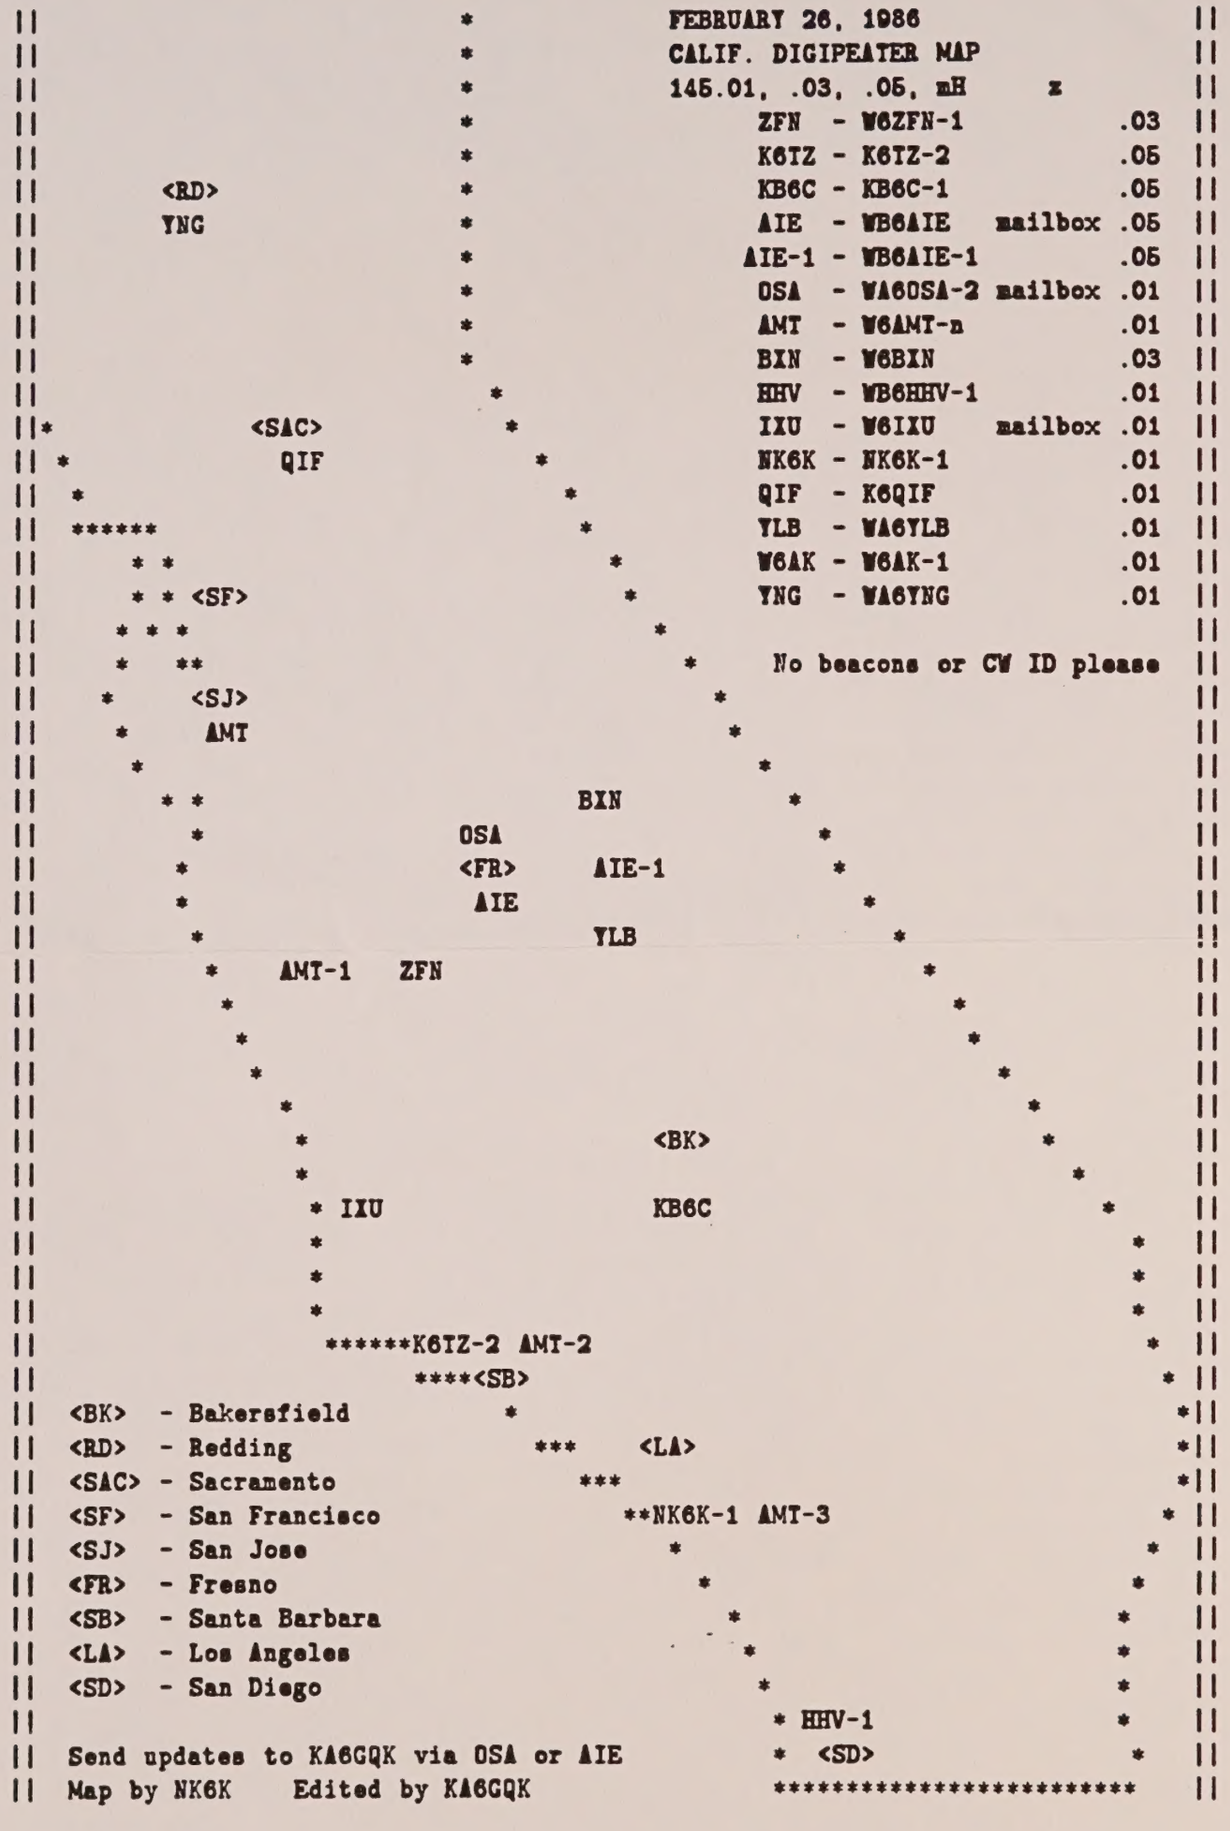
FEBRUARY 26, 1986
CALIF. DIGIPEATER MAP
145.01, .03, .05, mH z
ZFN - W6ZFN-1 .03
K6TZ - K6TZ-2 .05
KB6C - KB6C-1 .05
AIE - WB6AIE mailbox .05
AIE-1 - WB6AIE-1 .05
OSA - WA6OSA-2 mailbox .01
AMT - W6AMT-n .01
BIN - W6BIN .03
HHV - WB6HHV-1 .01
IXU - W6IXU mailbox .01
NK6K - NK6K-1 .01
QIF - K6QIF .01
YLB - WA6YLB .01
W6AK - W6AK-1 .01
YNG - WA6YNG .01
No beacons or CW ID please
<RD>
YNG
<SAC>
QIF
<SF>
<SJ>
AMT
BIN
OSA
<FR>
AIE-1
AIE
YLB
AMT-1
ZFN
<BK>
KB6C
IXU
K6TZ-2 AMT-2
<SB>
<LA>
**NK6K-1 AMT-3
HHV-1
<SD>
<BK> - Bakersfield
<RD> - Redding
<SAC> - Sacramento
<SF> - San Francisco
<SJ> - San Jose
<FR> - Fresno
<SB> - Santa Barbara
<LA> - Los Angeles
<SD> - San Diego
Send updates to KA6GQK via OSA or AIE
Map by NK6K Edited by KA6GQK

Yes, the change is probably inevitable. So is the sun's red-giant phase. Which will occur first (or be more disruptive) is an open question.

Phil

M 25241 Howard Goldstein (N2WX,2987) 12/22/85 9:03 AM L:14
KEYS:/FLORIDA THRASHES SAREX2 SOFTWARE/
TO: (Group 95)

On the evening of Saturday, December 21, Henry WA4HX% flew a TNC 2 with SAREX2 (release 1.1.2a), and a Motorola HT from South Florida up the middle of the state to near Lakeland and back down.

899 ROBOT QSOs were made, no real problems were reported, nor were any observed from my QTH.

However the combination of DIGIPEAT ON and FULLDUP ON was terrible. Whenever someone tried to digi thru it the TNC would kick on its transmitter right over acks, UAs, etc.

I suggest that the actual mission default ONE of these parameters to OFF. Comments?

73 Howie

PS I'd also make a recommendation that TNCs which send SABMs while connected should not be allowed, but we don't have to get into that...

C2974 CC141 Howard Goldstein (N2WX,2987) 12/21/85 9:20 PM L:7
KEYS:/ONLY $99 DOWN $99 A MONTH FOR 99 YEARS/
A: 140

I keep hearing that users are being horse traded into CCITT networks but there isn't any proof. I never hear that the people who sold them are on unemployment lines.

Where's the one iota to back up dem claims? How can a market survive, and grow, if customers are disatisfied and things are as bad as they're alledged...

C2974 CC142 Phil R. Karn (ka9q,2979) 12/22/85 3:26 AM L:33
(ORIG.) 12/22/85 2:53 AM L:25
KEYS:/CCITT/WITH FRIENDS LIKE THAT WHO NEEDS ENEMIES?/

I'm constantly reminded that there isn't necessarily a rational connection between the price of a service and the cost of providing it, but a look at the existing networks might be instructive.

If you look at Telenet's tariffs and figure how much it costs to ship files around, you will note that one of the new 9600 bps dialup modems plus good ol' AT&T will blow them away, any time of the day or night. At certain times much slower speeds (like 2400 baud) will break even. Uninet is even worse, since they charge by the character instead of by the packet.

All this is further support for my contention that the only thing an X.25 network is good for is remote terminal concentration. The problem is, the users are getting a lot more sophisticated than that. They want to interconnect their local area networks, which come in all shapes and sizes. Some are datagram (Ethernet CSMA/CD busses, Pronet token-passing rings), others are virtual circuit (your canonical "data PBX", although I don't think they deserve the term "local area network"). The CCITT is used to a world where they own everything, including the telephone on your desk, and they are uneasy with the idea of dealing with a proliferation of "customer premises networks" that they don't have any control over.

Since the CCITT has never worked in a heterogeneous world before, I have very little confidence in their ability to promulgate standards that work in such an environment. Certainly not when there's another organization (ARPA) which has already proven its ability to work in a true "internetwork" environment. No doubt the CCITT members would love to do away with all these "non-standard" LANS and protocols, and legislate X.25 right up to all your computers, and of course they'd be happy to rent you all the necessary equipment...

Phil

C2974 CC143 Howard Goldstein (N2WX,2987) 12/22/85 9:58 AM L:12
KEYS:/AND NOW/THE REST OF THE STORY/
A: 142

Yes comparing Telenet with AT&T is further proof of my point that datagramme networks are useful ONLY for transferring one file.

Be serious, no one has enough money to tie up an AT&T voice circuit all day. Telenet is the method of CHOICE for interconnecting many geographically dispersed LANS. Moreover, a large number of private networks (big stock brokers for instance) who have a real need for real systems use REAL CCITT protocols in the real world, and don't they excercise their choice over network "flavor"? They're apparantly paying real money and most are not overtly subsidized by the DoD/ARPA, or any other government agency.

Meaningless comparisons seem to abound when the datagramme crowd tries forcing analogies...

Good Day!

C2974 CC144 Phil R. Karn (ka9q,2979) 12/22/85 4:03 PM L:31
KEYS:/PDNS AND NETWORK INTERCONNECTION/

How is my Telenet vs AT&T comparison supposed to reflect badly on datagram networks? Telenet is internally virtual-circuit based, while AT&T is just a fixed bandwidth circuit switching network. My point was that Telenet and the other X.25 carriers are not as "cost effective" as many people believe, unless your bandwidth requirements average at very low values over long periods of time (i.e., represent interactive terminals). People interconnecting local area networks such as Ethernets want their computers to be able to talk to other computers, and in this situation you need BOTH fast "setup" time and the ability to ship large files economically. Perhaps the picture will change as more of the PDNs install fiber backbones and competition brings prices down. However, the LAN-interconnect users will still be saddled with the hairy X.25 VC interface originally designed for slow speed terminals when they're trying to gateway their 10 megabit Ethernets together. Something different is desperately needed, and ISDN certainly isn't it.

I think you'll find from reading just about any paper on the subject of heterogeneous network interconnection that the datagram approach is the ONLY workable way. Xerox is a good example of this; we haven't talked about it much, but the Xerox Network Systems set of protocols had common beginnings with the ARPA Internet protocols, and despite differences in detail (address field formats, the distinction between the Internet and Transport layers, etc) they both rely on the notion of a standard "Internet Packet" datagram which is encapsulated in whatever link level protocol happens to be available.

I could cite plenty of other networks which were all built around the datagram concept, dispelling the erroneous notion that only the DoD likes them, but I'd probably bore you with the length of the list.

Phil

C2974 CC145 Phil R. Karn (ka9q,2979) 12/23/85 6:41 AM L:160
KEYS:/A MORE EFFICIENT LINK LEVEL PROTOCOL/COMMENTS WELCOME/

Recently I've been thinking about ways to improve the performance of link level protocols on noisy links, and I came across a very interesting idea.

In the paper "ALOHA Packet Broadcasting - A Retrospect", by Binder, et. al., published in the 1975 (10 years ago!) National Computer Conference, are the following interesting sections:

[KA9Q Note: the "MENEHUNE was the ALOHANET's packet switch, analogous to our TNC. Its name is an obscure wordplay. The ARPANET people were calling their packet switches "IMPs" (Interface Message Processors), so the ALOHA people simply chose the Hawaiian word for "imp", as in "mischevious dwarf".]

"Error control for broadcast channel data packets (MENEHUNE to user nodes) involves some special considerations. For efficient operation, the usual positive acknowledgment scheme in which the ACK's themselves are not acknowledged depends on a high probability of the ACK's being successfully

received. However, an ACK sent from user nodes must compete with data traffic in the random access channel. At full channel loading each random access packet must be retransmitted an average of 1.7 times, which means each data packet or ACK must be sent a total of 2.7 times on the average before it is successfully received. But in order to force retransmission of the ACK's, the data packet being acknowledged must also be sent an average of 2.7 times by the MENEHUNE - even though it was received correctly the first time!....

[ka9q note: these numbers come from the fact that running a pure ALOHA channel at the optimum throughput level of 1/2e or 18.4% requires an "offered load" of 1/2 or 50%. I.e., the channel will have at LEAST one carrier on it 50% of the time, but only 18% of the time will there be ONLY one carrier. This means that each packet has a probability of (1/2e) / (1/2) = 1/e = 36.8% of making it across, and therefore packets must be sent on the average 1/(1/e) = e = 2.718 times to make it across once.]

"...to eliminate the unnecessary repetitions of data packets from the MENEHUNE, it is also necessary to acknowledge the ACK. That is, the ACK sent by a user node is timed out and retransmitted until an acknowledgement for it is received, just as for data packets. If another packet is waiting for transmission to the node at this time, its transmission with the next sequence number constitutes the ACK to the ACK; otherwise, a short ACK-ACK packet is sent by the MENEHUNE. This can be easily shown to result in significantly less total channel overhead, at the expense of more complication in the node implementation."

I was intrigued by this idea and decided to study it further. What follows is my initial analysis of the ACK-ACK protocol and how it compares to the current AX.25 protocol.

In AX.25, ACKs are not acknowledged. This means that if the probability of getting a packet across a given channel is p, then on the average, A = 1/p transmissions of the ACK by the receiver will be necessary for the sender to get one back. To elicit 1/p ACKs from the receiver, the sender has to send D = 1/p * A = 1/(p**2) data packets. For example, if p = .5 (50% probability of getting a packet across the channel), then you can expect, on the average, A = 2 transmissions of the ACK and D = 4 transmissions of the data packet before sender can go on to the next data packet. If p = .25, then A = 4 and D = 16.

Note I've assumed for sake of illustration that p is equal in both directions; in the "real world" a data packet is a bigger target than an ACK and is therefore even more likely to get clobbered.

To show the effect of ACK-ACKs, I'll define a number N, the ratio between the sender's data-retransmission timer to the receiver's ACK-retransmission timer. We assume that this number is always greater than 1. (If it wasn't, the sender would always time out and resend the data before the receiver ever gets a chance to retransmit its ACK). N therefore represents the number of times the receiver will attempt to retransmit its ACK before the sender unnecessarily resends its data packet (which is what we'd like to avoid).

The probability P of getting an ACK through at least once in N transmission attempts when each attempt has probability p of succeeding is

```
P = 1 - (1 - p)**N
```

Make N large enough, and P approaches 1 (i.e., the ACK will almost certainly get through.) Of course there will still be plenty of data packet retransmissions, simply because (1-p) of them will never make it to the receiver in the first place. So the total expected number of data packet transmissions D changes from 1/(p**2) in the present case to (1/p) * (1/P) with ACK-ACKs. Let's plug in some numbers and see what effect this has on D, the expected number of data packet transmissions.

```
p = .5, N = 5, no ACK-ACK case:
     D = A * 1/p = 1/(p**2) = 1/(.5*2) = 4
same, with ACK-ACKs:
     P = .969
     D = (1/p) * (1/P) = 2.06, a 49% decrease
```



```
p = .25, N = 5, no ACK-ACKs:
     D = 16
same, with ACK-ACKs:
 P = .763
     D = 5.24, a 67% decrease
```

You may be asking "what about the extra overhead of the ACK-ACK"? Recall that a new data packet can serve as an ACK-ACK. So long as there is a steady stream of traffic, "standalone" ACK-ACKs are never transmitted; the next data packet provides that function. Only at the end of a traffic burst is an extra packet generated, and this seems a small price to pay to save all those unnecessary retransmissions.

So clearly we have something here. The question is how to implement it.

I am not certain that this technique can be adapted easily to LAPB, the connection-oriented upper sublayer of AX.25 Level 2. It may be possible, but it would take some work. LAPB was originally intended for highly reliable point-to-point links, and radical surgery may be necessary. In any event, LAPB is an unnecessarily complex link level protocol for a datagram network, so I already have a motivation for doing a lobotomy instead.

Fortunately, there are "hooks" in AX.25 that allow us to dispense with LAPB and implement the ACK-ACK protocol while still conforming to the AX.25 addressing format. We need three types of packets to implement the ACK-ACK protocol: a data packet, a regular ACK packet (sent in response to the data packet) and an ACK-ACK packet (sent in response to the regular acknowledgment when no additional data is available). Here's one possible assignment:

```
data packet - UI frame with or without Poll bit set
regular ack packet - UA, with or without Final bit set
ACK-ACK - empty UI frame with Poll bit cleared
```

The poll/final bits allow either or both acknowledgments to be optional. For example, stations in a datagram network operating on high quality links with very low packet loss rates might elect to operate without the overhead of link level acknowledgments; in this case, each data packet is sent in a UI frame with the poll bit cleared. Similarly, when regular acknowledgments are requested (i.e., with the poll bit set in UI data frames), the receiver uses the final bit in the UA to indicate whether an ACK-ACK is requested. Since a new data packet can be sent in place of an ACK-ACK, making both out of UI frames simplifies the receiver implementation. In each case, setting the poll/final bit means that some sort of response is expected from the other side, which also simplfies the implementation.

Notice that there are no sequence numbers, since a sliding window protocol such as LAPB is unnecessary on a half duplex channel. The protocol operates in a lock-step fashion; there is always a one-for-one relationship between a data packet and its acknowledgement, and this simplifies things enormously. Since packets cannot "cross in the mail" on a shared half-duplex channel, there is never any ambiguity as to which data packet an ACK refers to (note that this assumption is violated and the protocol breaks down in two situations: a long digipeater string where the end stations cannot hear each other directly, and on a long-delay satellite channel where the round trip delay exceeds the packet transmission time. Digipeaters we're already getting rid of, and I'll worry about high speed satellite channels when we get one.)

The only "state" required in the protocol is a buffer for the most recently sent or received packet. (The transmitter may need to resend it and the receiver needs to detect and filter out duplicates transmissions). However, once a packet has been sent and fully acknowledged, no state needs to be maintained, a MAJOR advantage when servicing a large number of mostly idle local "links" with limited memory.

Comments and suggestions on this idea are welcome.

73,

Phil Karn, KA9Q

C2974 CC146 Jack Brindle (WA4FIB,2963) 12/27/85 12:42 AM L:73
KEYS:/LAYER 3 AND 4 BOTH NEEDED/PERHAPS TCP/IP HAS SOME GOOD IDEAS!/

I guess I will now jump into the discussion. Howie and I have been having a "private" discussion on the layer 3 - 4 issues, which should be carried on here.

First, we should all realize that ISO and CCITT are two completely different organizations with seperate goals and purposes. They just happened to have adopted similar communications models.

The first thing we must take care of in our system is to assure the delivery of data from one end user to another. This is accomplished in the OSI model with the transport layer. The network layer is incapable of doing this function since it does not "see" the information that is needed to perform the function. The network layer simply guides the data through the network from node to node. If you don't believe this then it is time to look at the information supplied by both ISO & CCITT.

Because of this, any implimentation of a network layer without the transport layer will suffer lots of problems. An example of this on a smaller level is the current state of file transfers on packet. File transfers are generally not done in packet radio (RLI BBS transfers are really message exchanges - look at the problems had in long file exchanges). The reason is the lack of a decent transport mechanism above the link layer. The link layer can "guarantee" the delivery of a small number of packets, but it cannot "guarantee" they will all arrive, and get there in order. Attempts at using XMODEM have tried to overcome this problem, but have met with other difficulties.

This is not an attempt to "advertise" my session layer protocol, but rather points out that node-to-node error checking is not good enough for the file transfers than will be required in the future. I maintain that packet radio in the future will be quite different than it is now, and must support far more and better services than we are willing to support now. The place for error checking is in the end user's systems. This is done at layer 4. Now if anyone wishes to release a layer 3 coding with a layer 4, be aware that ridicule will come. (But then it has come to those who wrote the layer 2 spec we all use...)

Probably the biggest advantage that the TCP/IP folks have to offer is that they have a "well-designed" system. Read that again - they have a system. One that was designed for all the pieces to fit together. One of the biggest problems in the packet world today is the lack of CCITT standards above layer 3 (at least until recently). Although packet switches work well at switching packets, it is extremely difficult to get terminal controllers from different manufacturers to work together. This is because each vendor used a different spec for layers 4, 5, 6 and 7. Not even the ISO and CCITT agree on the exact purpose of layer 7. (CCITT puts its actual drivers and hardware (read screen drivers) interface code here, whi while the ISO uses it as an interface inside a computer). This problem also exists in ham packet radio. We don't have a system design. We are building the system up from the inside out. What results will invariably become a kludge. The thing that really bothers me is that this "kludge" is quickly being adopted by the low-end players of the commercial world!

Now I propose that the arguing over end-to-end acking versus node-to-node be placed aside. The answer, of course, is that we really need both. Instead why don't we draft a layer 4 spec and get started on working to design and code it.

Of course, we could follow Phil's lead and adopt TCP/IP with all of its ramifications including a complete system design...

Phil; I am indeed interested in running your code. I have just finished up MacPacket/TNC2term (and shipped my backlog last Monday, the 23rd), and will shortly be moving from temporary quarters into a house. I will be able to devote time to integrating the code into my MacPacket package for demonstration purposes. I managed to lose your address, please give it again. This time I promise to get a large envelope to you (with suitable postage, whatever that is) so that I can get a diskette and whatever docs you would like to send.

I again also extend the offer to the VC folks. Give me a good layer 3 & 4 implimentation and I will code it into MacPacket also. I would love to be able to run this stuff at Orlando to help out in our discussions.

Another note: Those of you on Compuserve, take a look at the files FCC1.DOC and FCC2.DOC in Public Access. (GO ACCESS). They contain a very interesting document from a ham in Calif that is proposing a Public Data Service to run on a portion of the "unused" amateir 6 meter band at up to 1 MB/s! I'm not sure if this guy is just dreaming, but the FCC has decided to take some sort of action on his petition. This looks like something that we need to channel back into ham radio. I'm not a supporter of "no-code", but perhaps it is an idea whose time has come!?

Good night all - I hope everyone had a Merry Christmas.

C2974 CC147 Phil R. Karn (ka9q,2979) 12/28/85 3:34 AM L:53
KEYS:/NOW THAT WE HAVE LEVELS 3 AND 4 LETS MOVE ON/

Jack, thank you very much for your comments. I am happy to hear them; I've been saying similar things for over a year now!

My TCP, as a "layer 4" protocol, is up and running. It implements the full-blown ARPA TCP spec as described in RFC-793 and MIL-STD-1778 with two exceptions: security (which I assume we have no use for) and urgent data (which I can't figure out how to use, although it wouldn't be hard to add). Other than that, there was no reason to change the spec to suit our needs, with the result that my implementation is fully compatible (and has been tested with) the various other implementations out on the Internet. I hereby toss it out on the table for consideration as our end-to-end Transport protocol. There is no reason to "roll our own", thus adding to the layer 4-5-6-7 confusion you alluded to. In response to comments regarding the programming interface, I've reworked it to make the facility more general; there are now optional "upcalls" (also known as "pseudo interrupts" for transmit buffer availability and connection state change as well as receive data availability.

As I mentioned earlier, the ARPA Internet has no need for a distinct session layer (level 5); its few functions are assumed by the "well known port" concept in TCP (and UDP). As for presentation layers, this depends on the application. Since remote login is likely to be an important initial application, I'd like to suggest that we consider the ARPA TELNET protocol. This is a simple, extensible technique where various options (e.g., which end of a connection is to perform character echoing) can be negotiated. For example, a user connecting into a UNIX system (which conventionally operates in "full duplex" mode with per-character transmissions and remote echoing, to allow use of fancy command line and screen text editors) will receive a WILL ECHO offer from the host. Your Telnet, which is under your control, can respond with either a DO ECHO or a DON'T ECHO command, which either accepts or rejects the remote host's offer. Many other options exist, including ones to control the faucets on the kitchen sink, but only a small set (about 5) are actually recommended for use and only one or two (echo being the most important) are actually used widely. The mechanism is there, however, for other options to be implemented, and their use can be negotiated with other systems as needed.

Other presentation level protocols of interest include the FTP (File Transfer Protocol) and the SMTP (Simple Mail Transfer Protocol) with self-evident uses. SMTP really is simple to implement, as there are only about 5 commands. FTP is quite a bit more complicated, since it is one of the oldest commands in the ARPANET (long predating TCP/IP and about as old as TELNET), and although it is widespread there may be simpler alternatives. I would like people to look at SMTP, however, since it represents a much "cleaner" way to transfer mail messages than the current highly ad-hoc of linked WØRLI bulletin boards. I can provide documents on these protocols, along with sample sessions, if there is interest.

Phil

C2974 CC148 Scott Loftesness (W3VS,2961) 12/29/85 3:27 AM L:10
KEYS:/BRINDLE GET WITH IT/

Hey, Brindle, when are you going to share the specifications on those new protocols of yours? I thought you told us over a month ago that they're be here Real Soon Now! Did I miss 'em?

By the way, Stoner's proposal for the Public Digital Radio Service sounds good to me. Hell, implementation of that service and the definition of the protocol should keep you protocol warriors fully employed for years!

[Flame off!]

M 27278 Howard Goldstein (N2WX,2987) 12/31/85 2:59 PM L:41
KEYS:/AND BEST WISHES FOR A PEACEFUL AND PROSPEROUS NEW YEAR/
TO: (Group 95)

to: All
fm: You guessed it, Howie
re: It's 10PM, do you know where your vaporware is?
dist: Wide and far

The local digipeater (N2WX-7 a.k.a MLB) is now also a network node. I am not sure if it is the first instance of ax.25 13 AND the stuff to do the pad functions for folks unlucky enough not to have a TNC 2, but I include the example below of a short session with the Orlando mailbox using the PAD.

Networking, and ax.25 level 3 are REAL. They're HERE. ANYONE can use them.

Of course this will all be available for anyone's use from the Orlando conference.

---

actual unretouched example of the WORKING ON THE AIR pad:

```
:\
cmd:c n2wx-7
cmd:*** CONNECTED to N2WX-7
{4
gator 2 pad 03100305724 port A

enter: call [,digi1 [,digi2 [,digi3] ] ]
to?k4aho orl
*** pad: connection reset
K4AHO Packet BBS, Orlando, FL, new msg's 1126 to 1125
Type X for long menu, H for Help
Last Msg 1125, 59 Active Msgs (B,D,G,H,I,J,L,R,S,T,U,W,X) >
Oh, Hello Howie
Last Msg 1125, 59 Active Msgs (B,D,G,H,I,J,L,R,S,T,U,W,X) >
l1 3
Msg# TR  Size  To    From   @ BBS  Date     Title
1125  N   248 WD4KAV WB4TRJ        851231 greetings
1120  N   208 ALL    WD4IXD        851230 voice net pt.2
1119  Y   257 K4AHO  N4LXI         851230 BBS
Last Msg 1125, 59 Active Msgs (B,D,G,H,I,J,L,R,S,T,U,W,X) >
b
*** DISCONNECTED
N2WX-7>K4AHO,ORL <UA>
N2WX-7>N2WX-1 <D>
```

M 27337 J. Gordon Beattie Jr. (N2DSY,2990) 12/31/85 8:33 PM L:29
KEYS:/WE DID IT IN '85 !/
TO: (Group 95)

I have been a bit quiet for the last few months frankly because I hate going around in circles, but I will make a comment on the end of the year happenings.
(Pardon the condescension, but some of the comments have been frustrating.)

We have gone a year and we finally have a few working implementations of level 3 network components using the OSI and ARPA models.
THIS IS GOOD !

Regarding requests for protocol specs: we should stick to readily available standards in order that we are able to scoff what we can in the way of experience, software and hardware. I am a bit concerned that Jack's work is undocumented, but I am more concerned that it is an unknown protocol with actual work done on an implementation. Vacuums are a horrible place to live or work !
The spec is important Jack, but the work your doing with the terminal programmes needs to have something that others can talk to.

Last but NOT LEAST: THANKS TO LYLE, PETE, AND HOWIE
for their work on the TNC-2 and NNC.
WE HAVE LEVEL 3 VIRTUAL CIRCUIT NETWORKING IN 1985 !
Now WE must go and install the networks !
Have a GOOD YEAR !

73, Gordon, N2DSY

P.S. C U in Orlando !

M 27386 Pete Eaton (WB9FLW,2970) 1/ 1/86 1:47 AM L:11
KEYS:/PAT YOURSELVES ON THE BACK-NOW FOR A SIMPLE PROJECT/
TO: (Group 95)

Gordon:

Thanks for the kind words on TNC-2. Everyone on DRNET, and many others who are not, also sweated the last year out. So everybody pat yourself on the back!

Ok, now for a "simple project", anybody want to build a RF deck?

Happy New Year!

Pete

M 27613 GARY GARRIOTT (VITA,568) 1/ 2/86 8:48 AM L:6
KEYS:/TNC-2 MATING/
TO: (Group 95), VITA

DOES ANYBODY OUT THERE KNOW OF ANY MATING PROBLEMS OF THE TNC-2 WITH RADIO SHACK MODEL 100 OR 200? HAVE HEARD RUMOR TO THIS EFFECT... WE'D LIKE TO TAKE EITHER THE 100 OR 200 TO ETHIOPIA...

GARY

M 27638 Thomas A. Moulton (W2VY,995) 1/ 2/86 10:02 AM L:3
KEYS:/TNC-2 MATING/
A: 27613 TO: (Group 95), VITA

One problem I have seen is that when using the built in terminal program the M 100 seems to send X-OFF and never sends an X-ON to allow the tnc to send to the M100 again

M 27809 Pete Eaton (WB9FLW,2970) 1/ 2/86 7:43 PM L:13
KEYS:/THIS IS GOING TO BE HARD TO BELIEVE/
TO: (Group 95)

To: All
From: Pete WB9FLW
Subject: RLI Mailbox/Gateway On-Line
Dist: Open

I know this is going to be hard to believe but as of 4:30 P.M. Central Time today I finally got the RLI system back up including the 20 meter Gateway. I'll leave the system on 24 Hrs aday, be looking forward to chatting with you folks on the "other" network!

C2974 CC150 Phil R. Karn (ka9q,2979) 1/ 4/86 12:04 AM L:19
KEYS:/NETWORK PRE-EMPTION/

As I mentioned in my writeup on the Internet protocols, I support the notion of packet precedences or priorities. It is possible to queue packets for transmission based on these priorities; for example, a node would transmit all packets marked PRIORITY before transmitting any packets marked WELFARE, and so on.

It is also possible for a node to preempt (i.e., abort) the transmission of a packet should a higher precedence packet need the same link. There are arguments each way; clearly preemption will grant the lowest possible delay to urgent traffic when the links are already busy with routine traffic, but at the expense of lower overall efficiency (because of all the aborted packets). I am looking for suggestions on the best way to handle this. I could preempt whenever a higher priority packet needs the link; I could preempt only for the highest priority (EMERGENCY), or I could preempt if the relative priorities differ by more than 1 level. Comments, especially from traffic handlers, would be appreciated.

73, Phil

C2974 CC151 Lyle Johnson (WA7GXD,2973) 1/ 4/86 12:36 AM L:36
KEYS:/NNC PROGRESS REPORT/

To: Anyone Interested
Fm: TAPR
Re: NNC Prototype debugging progress
(03 jan 1986)

The NNC prototype testing is progressing fairly smoothly. The test unit is chugging away at a system clock of 4.608 MHz (due to CMOS Z80"A" peripherals being available, while "B:" parts are not).

Using test code provided by N2WX, I have verified the operation of one channel of 1 SIO running HDLC data in a digital loopback mode at a data rate of 76.8 kbps (it can't hack 153.6 kbps in fulldup with the test code, however...). This is polling, not interrupt-driven. I have also verified that I can talk to both channels of both SIO/2s.

Similarly, I have verified that I can talk to both channels of the PIO and ran a simple test program tonight to drive a "Centronics" (well, OkiData) printer.

I have also written data to and read data from the SCSI chip (5380) but it is basically gibberish, I haven't attempted to make it do anything useful.

Finally, I have patched the BIOS to support the 4.608 Mhz clock rate and all the software I have written/used is running on the NNC prototype under the Z-system.

I have a pair of Mitsubishi drives on qualifaction loan, and Fujitsu and Tandon drives are enroute for the same purpose. The Mitsubishi drives are quiet and fast...

Oh yes, the prototype is using 8k-byte SRAMs - NEC didn't accept my offer to purchase 32k-byte parts at prices I was willing to pay! (but, I did get 1 free sample!!!).

More later. Lyle

C2987 CC246 Lyle Johnson (WA7GXD,2973) 1/ 5/86 12:54 AM L:10
KEYS:/LOTSA RAM INNA CHIP/WORKS FINE/

I have successfully interface a 32k byte SRAM to a tnc 2 rev 2 board. Simply remove bb power (as well as regular power!), remove the two 62641s from u24 and u25, and install the 43256 SRAM at u25. reconnect the bbram link, power up, initialize and go for it.

Oh yes. the 43256s cost only about $100 each in small quantities...

Howie, got any code that can make use of 32k bytes of bbRAM?

Lyle

M 29254 GARY GARRIOTT (VITA,568) 1/ 7/86 7:48 AM L:21
KEYS:/URGENT FEEDBACK NEEDED/
TO: (Group 95), VITA

YESTERDAY I HAD A CHAT WITH DON MERTEN (K2AAA) WHO WAS REFERRED TO ME BY ALEX MONTARE, KA2VLP. DON HAS WORKED FOR MANY YEARS WITH DIGITAL TRANSMISSION METHODS OVERSEAS AND INSISTS THAT AMTOR, RATHER THAN PACKET, IS THE PREFERRED MODE FOR HF. ACCORDING TO DON, THIS IS SO BECAUSE FREQUENCY STABILITY IS MUCH MORE DIFFICULT ON HF PLUS THE RELIABILITY FACTOR, IE, FEWER BYTES PER BURST ON AMTOR TRANSLATE INTO MUCH HIGHER EFFICIENCY GIVEN STATIC CRASHES, ETC. DON IS COMING TO TOWN TODAY AND I MAY SEE AN AMTOR DEMO TONIGHT.

NATURALLY, THIS IS SOMEWHAT DISTURBING SINCE WE HAD ALREADY PURCHASED THREE TNC 2'S FOR THE ETHIOPIA PROJECT, BUT DON HAS ME JUST ABOUT CONVINCED THAT WE BETTER GO OVER THERE WITH AN AMTOR CAPABILITY AS WELL AS PACKET WHICH MEANS THE AEA "PACK-RAT" (PK-64)/COMMODORE 64 COMBINATION.

GUYS, IF WE GO TO ETHIOPIA AND IT DOESN'T WORK, OUR (I SHOULD SAY MY) GOOSE IS COOKED. I REALIZE THERE ISN'T MUCH COLLECTIVE EXPERIENCE OF PACKET ON HF IN THIS GROUP, BUT WOULD APPRECIATE SOME FEEDBACK ON THIS NEW DEVELOPMENT.

COULD DON BE RIGHT?

GARY

M 29258 J. Gordon Beattie Jr. (N2DSY,2990) 1/ 7/86 8:15 AM L:16
KEYS:/YOUR GOOSE IS NOT COOKED/
TO: VITA, (Group 95)

Gary,
There are those who can make a case for AMTOR/SITOR, but there is more to the story:

packet will NEVER allow a character error of any kind - AMTOR will.
there are bunches of packet stations of 14.103 that ARE working fine.
There is a 10 MHz frequency also and I can't recall it.
Packet uses a Timeout on error then retransmit philosophy, AMTOR/SITOR uses a FEC technique. Packet uses a multipacket rotating window acknowledgement system and AMTOR/SITOR uses a basic ack for each data system.
Most HF packeteers keep their packet lengths short so that the packets are less susceptable to hits during transmission. This takes care of the reliability question.
73, Gordon N2DSY

M 29322 Thomas A. Moulton (W2VY,995) 1/ 7/86 12:42 PM L:9
KEYS:/URGENT FEEDBACK NEEDED/
A: 29254 TO: (Group 95), VITA

I think that Don is assuming that you will use 128 byte packets, with packets that long he is right.

I don't know what the guys on HF mail forwarding use, but I think you could use packet sizes of 64 or 32 bytes long.

I have sent many messages that have gone via HF, it DOES work!
I realize that it is about twice as far to ET
Lets see what the rest of the group has to say.

M 29692 GARY GARRIOTT (VITA,568) 1/ 8/86 7:33 AM L:10
KEYS:/THANX FOR FEEDBACK/
TO: (Group 95), VITA

THANX EVERYONE FOR THE COMMENTS...I'VE NOW SEEN DON'S SET-UP AND IT IS IMPRESSIVE...THE FACT THAT HE HELPED INSTALL THE VERY RADIOS IN ETHIOPIA SOME YEARS AGO IS ANOTHER FACT...HE DOESN'T THINK TIMING WILL BE A PROBLEM VIS-AVIS THE X-R RELAY....AM LEANING TOWARD HAVING A DUAL DEMO CAPABILITY JUST TO BE ON THE SAFE SIDE...

FURTHER THOUGHTS WELCOME...

GARY

M 29707 Paul Newland (ad7i,2978) 1/ 8/86 8:39 AM L:53
KEYS:/AMTOR/PACKET/HF/
TO: (Group 95)

FROM: paul newland, ad7i
TO: gary garriott
SUBJ: packet/amtor on hf
DIST: limited
DATE: 86.01.08

gary -- allow me to add my comments to those already made.
I hope that it doesn't further confuse the issue.
I have a lot of HF amtor experience but NO packet experience at HF. So, I come with some strong biases.
Statements have been made that packet never allows an error to get through. I would recommend that the author of such statements rephrase that to be less misleading. I would suggest that packet offers extremely good error control, but not perfect. All single bit errors will be detected but there is a non-zero probability that a packet with multiple errors might be accepted as correct. However, the probability is remote (~1E-5).
The MODE-A operation of amtor alluded to is certainly not Forward Error Control (FEC). It is an ACK/NAK as opposed to the ACK/NOTHING protocol of packet. one strong advantage of packet here is that the ACKs themselves of packet are much better protected against errors than the ACKs of amtor.

Several folks have attempted to use fairly unsophisticated analytic methods to compare amtor and packet. They have resulted in unsophisticated results. I wouldn't place too much faith in such results. I can tell you this however, both work and work very well. I will speculate (and I underscore the word "speculate") that amtor will work better in a slugfest on 80 or 40 meters when packet won't get through, because the channel error rate is 1E-2. The price you pay such poor channel performance is no data throughput with packet and a low throughput with poorer error probability for amtor. Look at Peter Martinez's article in QST 1981, pp25-27. If I remember right, he gives some throughput and error probability numbers for different channel error rates for amtor.
I don't recall your exact application but if you want to move a large amount of computer data at 10 MHz or higher and can wait for up to 24 hours (i.e., wait for best propagation of the day), and you insist on an error rate of 1E-9, then packet is probably the best way to go. It's tough to argue with success; packet works every day on 10 and 14 MHz.

If you want to move a smaller amount of data (i.e., 3 pages or less) and want to do at on 80 meters and do it right now and you can live with an error rate of 1E-5, then AMTOR will suit your needs best.
Bottom line, both work fine! I doubt that your goose will even need to wiggle its wings, let alone spending time in the pot.

73, paul

M 29768 J. Gordon Beattie Jr. (N2DSY,2990) 1/ 8/86 12:24 PM L:9
KEYS:/AMTOR/PACKET/HF/OOOOOPS !/
TO: (Group 95)

I agree with Paul, but I'm not sure if I'd want 1E-5 when I could have 1E-9. Sometime a promptness is more important though...
In reference to the term "errors", I would like to note that when a link breaks down (loss, resets, etc.) I want an indication. AX.25 does provide that to the user. I don't know what AMTOR does, but I would assume it has a recovery mechanism.
Pardon the faux pas Paul.
73, Gordon

M 11 Scott Loftesness (W3VS,2961) 1/ 9/86 1:12 AM L:119
KEYS:/HAPN PC-PACKET BOARD/
TO: (Group 95)

Msg#	TR	Size	To	From	@ BBS	Date	Title
231	BY	5291	ALL	AA4RE		860106	IBM PC based TNC announced

****
**
** Hamilton Area Packet Network introduces the **
**
** HAPN PC-PACKET ADAPTER **
**
****

The HAPN adapter is unique in terms of its versatility and ease of operation. It is designed for the IBM PC and compatibles. The IBM PC or compatible is rapidly becoming a leader in the HAM-world due its attractive price and excellent performance.

The HAPN packet adapter is a card that plugs directly into one of the slots of the PC. It contains a built-in modem which interfaces to the radio with a DB9 connector. The card is 8.5 inches (22 cm) long and contains a prototype area. The software is powerful and is a delight to use. The transfer of ascii or binary files is very simple by using the function keys on the keyboard. A popup menu allows you to set up the repeater routing table. Optional software is available for bulletin board operation. Experimental protocols such as VADCG V1 and V2 are also available.

Features of the H.A.P.N. packet adapter :

1. In addition to the American AX.25 protocol it also runs the Canadian experimental protocols V1, V2 as designed by Doug Lockhart (VE7APU) of the VADCG (Vancouver Amateur Digital Communications Group). (Doug is a world leader in modern packet radio since 1979 and founder of the VADCG.)

2. Excellent software. The software provided by HAPN is user friendly, easy to operate and very powerful. Available at present:
   a. AX.25 host support (included)
   b. bulletin board program PCRBBS (optional)
   c. file transfer programs (optional)
   d. experimental protocol support for VADCG V1, V2 (optional)

3. The HAPN adapter plugs into a PC, PC-XT, PC-AT or compatible. No external cabinet or power supply is required.

4. Simple to interface to your radio. It requires RX audio, TX audio and push-to-talk. The carrier detect (radio channel busy) is generated internally from the received data or optionally brought out from your squelch circuit.

5. The built-in modem is Bell 202 compatible. It uses the XR2206 and XR2211 chipset. The tones are 1200 and 2200hz. The data rate is 1200 baud. A hardware watchdog circuit prevents possible TX runaway problems as a result of power glitches or lightning.

   NOTE: The baudrate for the adapter is selectable by jumper from 75 baud to 9600 baud. (The on-board modem supports up to 1200 baud.)

6. The board contains a prototyping area for personal use. The area is pre-drilled.

7. The adapter is based on the INTEL 8273 HDLC/SDLC protocol controller chip with a straightforward interface to the PC.

8. The low level software (device driver) is interrupt driven and stays resident in the PC allowing the adapter to function in the background. An application program such as bulletin board, file transfer or terminal program does not have to be resident for the adapter to function. In this mode, the user is alerted to a caller by a connect alarm utilizing the PC's speaker.

9. Easy customizing for the user. The user configuration such as callsign/SSID, auto logging of packet activity to disk, default repeater, connect alarm etc. are set up using the configuration screen and remembered by the program. These options can be changed at any time.

10. NOTE FOR PROGRAMMERS
    - Since the adapter uses the native code of the host computer no cross compilers are needed for anyone wanting to develop his own code.

ORDERING INFORMATION AND PRICE :

All prices are in US$ and include postage and handling in USA or Canada. Please add 3$ for overseas orders.

1. Assembled and tested board with AX.25 host software and self-test program (4 programs and documentation)	$199.00
2. Bare board with AX.25 software and construction details	$ 75.00
3. AX.25 software only	$ 40.00
4. Bulletin board & file transfer programs	$ 25.00
5. Advanced VADCG experimental software	$.25.00

Make checks and money orders payable to H.A.P.N. The address is :

HAPN
Box 4466, Station D,
Hamilton, Ontario,
Canada, L8V4S7

A Note on HAPN

The HAPN (Hamilton Area Packet Network) is a Packet Radio club founded in 1980 and dedicated to furthering the state of packet radio. It is currently involved in the design of a station node controller. The station node controller is capable of unifying a packet radio network by providing message routing, gateways and high-speed links with other nodes. Money raised with projects like this helps us pay for research and development.

HAPN is also involved in experimentation with 4800 baud modems, and experimentation with the newly developed VADCG protocol V3.

C2974 CC153 Harold Price (NK6K,2972) 1/ 6/86 3:59 AM L:393
KEYS:/LINKING SCHEME FORWARDED FROM COMPUSERVE/

Don't think I saw this on DRNET. It is forwarded from COMPUSERVE.
Item 23 entered Fri, Nov 22, 1985 (21:06) by Michael McClary
Proposal for a geometric organization of a ham long-haul network.

Things I've seen in previous items give me the impression that Michigan hams, at least, are thinking of hooking up the long-haul portions of a packet network with radio links pointed toward each of the four compass points, using the four ports on a Network Node Controller from TAPR or the like. This strikes me as being a bad fit to the real world, and I have a counterproposal. (It makes very good use of the NNC's four ports, by the way.)

If a ham long-haul network is set up, it will likely have to consist of network nodes which connect to high-speed long-distance links, plus a local area network (to get the individual hams' data on and off the long-haul links). Since it already is an expensive, multi-antenna and multi-radio box, it would be reasonable to place it in a location central to a large community, and make its LAN link strong and sensitive, making its coverage look much like a digipeater with the ability to "autopatch" into the long-haul network.

Each NNC's network-access LAN will be directly (i.e., no digipeater hops necessary) accessible from an essentially circular service area. In order to provide high-quality (no digipeaters between you and the NNC) service to as many hams as possible, these circles should cover the entire land area serviced. In order to avoid wasting resources, they should overlap as little as possible. This packing of circles is a hexagonal (honeycomb) array, which tiles the plane quite nicely. (Failing this, each should serve an area that has many hams, and the dead spots should fall where hams are few.)

The hexagonal array has a hidden benefit. The economic forces which governed the placement of cities, towns, and villages caused them to be laid out in a hexagonal array as well. (Actually, they form nested hexagonal arrays on varying scales - a few big cities in an array with large cells, thousands of villages on an array with tiny ones, and some of the village sites occupied by cities.) (You may not notice the hexagonal layout on a map because it is slightly distorted. The towns are spaced for equal travel difficulty, rather than equal distance, so they are a little closer together if separated by bad terrain.) As a result of this layout, you have cities right where you want to put your NNCs.

At first glance, it looks like each NNC will have 6 nearest neighbors, and thus will need 6 high-speed links. It turns out you can get away with three. Consider the following array:

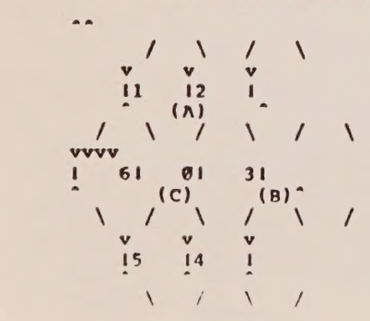

Suppose the numbered sites are major cities. Zero has 6 nearest neighbor cities. To talk to them directly would require 6 links, and each link would have to be able to carry over a distance sqrt(3) times as long as the reach of the LAN's transceiver. The intersections marked with "v", "^", "(A)", "(B)", and "(C)", however, would normally contain small towns, and would be just on the edge of the NNC's LAN range. Suppose each of these put up their own NNC, also with three long-distance links, but not necessarily with a LAN of its own. (Everyone in the town should be able to reach the LAN of one or more neighboring cities.) "0" would now talk to "A", "B", and "C", "A" would talk to "0", "1", and "2", and so on.

This sort of network can be built with four-port NNCs, and with four transceivers at the major nodes, and three at the minor nodes. It doesn't require extra-long range transceivers for the inter-NNC hops. (The directional antenna will give enough extra gain to make the inter-NNC link solid, even though it might be carrying a higher data rate.) REALLY large cities (which tend to form another, larger hexagonal structure superimposed on this one) can add longer-distance links, using 7-port NNCs (or 8-porters, with two LANs, which they'd probably need anyhow).The regular periodic structure and even spacing simplifies calculations of channel re-use and routing. In particular, there are easy-to-calculate routes around dead nodes, so it can be re-configured on-the-fly by reasonably simple programs. Finally, the cost of the individual network nodes is proportional to the size of the community they serve, and thus the cost is spread evenly over the community members.

How about it?Sound interesting?

RESPONSES

#1 Fred Steiner:
Question, Mike: Would each RF link be on the same frequency, operating half-duplex, or each on separate frequencies? What I'm getting at is how essential to network efficiency as you propose it is use of separate RF channels, and how important is Full vs Half duplex operation? When I say "same channel," I mean each link transmitter and receiver on the same frequency, so each is limited by the activity of the others in the Half-duplex mode. Would this effectively do in your system, especially if file transfers, or other data intensive activities were going on between two parts of the honeycomb or could the system tolerate such if a single frequency is used for the links. If a single frequency is too inefficient, how many different ones would be needed for optimum efficiency?
wa9hdd

#2 Michael McClary:
Let's see if I've got this correctly. You're asking:

1) Would using the same frequency for transmit and receive on one long-haul link degrade the performance of this geometry?

2) Would using the same frequency for all outbound, and perhaps a common frequency for all inbound, long-haul links degrade the performance?

3) What frequency allocations would be required for the optimum operation of this scheme?

Numbers one and two are not really matters of geometry. They are matters of the technology of the individual links. You can't share one frequency half-duplex for a long-haul link if you're going to operate all three links to a node in the same band. This is because the timing of the incoming messages on the other links is not under your control, and any that come in while you're transmitting on that end of the band will be lost to receiver quieting caused by your outgoing signal. Thus the incoming and outgoing signals must be a channel-pair, on opposite ends of the band, and this means that you may as well run them both continuously.

With regard to sharing the frequency between long-haul links to different sites, there are several ways this could be accomplished, with varying impact on throughput.

The simplest approach would be to allocate three channel-pairs for the long-haul links (N-S, SE-NW, and SW-NE), and re-use them over and over. You might also want to allocate another three pair for the very-long-haul bypass. (A third set might be useful for cases where the next hunk of the net isn't far enough over-the-horizon for isolation, or where there is some local interference, but this shouldn't really be necessary, because these are fixed stations, and you can adjust the transmitter output to make the capture of the long-haul link receivers solid.)
With this simple approach you only need three pair for the basic grid. If you look at the previous item, you'll find that at least four pair are already allocated for the square-layout net.

You could use one outbound frequency, and time-multiplex it between the three neighboring sites. This would reduce your output capability by a factor of three. Performance would probably degrade slightly less than that, because your load would not normally be balanced between the three directions. There would still be the problem, though, that you would have three times as much potential for incoming data as for outgoing, and this would degrade the usage of the incoming frequencies by this same factor of three (which shouldn't hurt, because they'd be talking to three nodes, too). (In a network like this, traffic THROUGH your node will normally be so much greater than local traffic TO/FROM your node that the latter can be neglected.)

Remember, though, that unlike the LAN, the long-haul links are between single fixed stations which are line-of-sight (just barely) and separated by (roughly) 120 degrees. This means you can use fixed, directional antennas. An antenna with two very sharp nulls at 120 and 240 degrees from the major lobe is almost trivial to construct: two driven elements, correctly spaced and phased. Add a third (forming a triangle) and the correct phasing harness, and you have three such antennas in one unit. (Also, with three elements, you can move the angles of the two nodes independently, to compensate for the real-world differences in angle between the cities.) Using such an antenna (and perhaps transmitters sharing a master oscilator) you can use the same frequency pair simultaneously in all three directions, build the network with only one set of long-haul frequencies, and experience no degradation.

The LAN you use for accessing your local node should be at the same end of the band on which it is transmitting (so it doesn't quiet its long-distance receivers) or perhaps in the middle of the band. You don't want to have adjacent cities using the same frequency, because this would cause trouble in the fringes, where stations would hear and be heard by two or three network nodes. This requires three frequencies, placed like this:

```
1     2     3  12     3     1

   3 1     2     3   1 2

1     2     3  12     3     1

   3 1     2     3   1 2
```

Another set of three frequencies, at the other end of the band, could be used if the intermediate nodes also were to establish LANs, producing overlapping coverage.

```
1     2     3  12     3     1
   c a     b     c    a b
   3 1     2     3    1 2
a     b     c  ab     c     a
1     2     3  12     3     1
   c a     b     c    a b
   3 1     2     3    1 2
```

Placing the single-frequency LAN at the same end of the band as the network node's transmitters means that the node's LAN receiver is quieted, and this reduces its range. It might be better to place it at the other end, and raise the strength of the adjacent node's transmitters, to compensate for quieting in the long-haul net. (I think the original approach is better, though, because it will tend to push local activity away from the band end where it would interfere with the long-distance links. It also tends to restrict the LAN to the network node's vicinity, (because the network node will shout down stations it doesn't hear, and they'll eventually give up and move to another frequency) and prevents interference by stations working the LANs at more distant nodes.)